WHAT IS CLAIMED IS:

1. A time:space:time switch fabric, comprising:

(a) a plurality of spatially distributed data switches, each one of said data switches further comprising:

(i) a first plurality of ingress ports;

(ii) a plurality of egress ports equal to said first plurality;

(iii) a space switch for selectably spatially rearranging data transmission between any one of said ingress ports and any one of said egress ports;

(iv) an ingress time switch coupled to an input of said space switch, said ingress time switch for selectably temporally rearranging data received on said ingress ports;

(v) an egress time switch coupled to an output of said space switch, said egress time switch for selectably temporally rearranging data output by said space switch;

(b) a plurality of spatially distributed data serializers, each one of said data serializers further comprising:

(i) an input bus for receiving signals to be routed through said switch fabric;

(ii) an output bus for outputting signals routed through said switch fabric;

(iii) a second plurality of egress ports, each one of said data serializer egress ports further comprising a high speed serial link selectably connectible to any one of said data switch ingress ports;

(iv) a plurality of ingress ports equal to said second plurality, each one of said data serializer ingress ports further comprising a high speed serial link selectably

connectible to any one of said data switch egress ports;

wherein:

(i) said data switch ingress and egress ports and said data serializer ingress and egress ports further comprise a composite fabric of said data switches characterized by at least:

(1) p planes, where p is a power-of-two integer less than or equal to said second plurality of said data serializer ingress and egress ports;

(2) s stages, where s is an odd integer number;

(3) a depth d, where d is a power-of-two integer less than or equal to said first plurality of said data switch ingress and egress ports;

said composite fabric comprising $p * s * d$ inter-connected ones of said data switches;

(ii) said data switch ingress ports of a first one of said stages are coupled to said data serializer egress ports;

(iii) said data switch egress ports of an s^{th} one of said stages are coupled to said data serializer ingress ports; and,

(iv) said data switch egress ports of each n^{th} one of said stages are coupled to data switch ingress ports of each $n+1^{th}$ one of said stages, where $1<n<s$.

2. A switch fabric as defined in claim 1, wherein:

(a) each one of said data switches further comprises:

(i) an active switch control memory for containing indicia defining currently active connections between selected ones of said ingress and egress ports for said one of said data switches;

(ii) a standby switch control memory for containing indicia defining next active connections between selected ones of said ingress and egress ports for said one of said data switches;

(b) each one of said data serializers further comprises:

(i) an active switch control memory for containing indicia defining currently active connections between selected ones of said ingress and egress ports for said one of said data serializers; and,

(ii) a standby switch control memory for containing indicia defining next active connections between selected ones of said ingress and egress ports for said one of said data serializers;

said switch fabric further comprising a controller for switching control of each one of said data switches and each one of said data serializers between said active and standby switch control memories of said respective data switches and said data serializers.

3. A switch fabric as defined in claim 2, wherein:

(a) each one of said data switches further comprises a start-of-frame delay register for containing a start-of-frame time signal unique to said one of said data switches;

(b) each one of said data serializers further comprises a start-of-frame delay register for containing a start-of-frame time signal unique to said one of said data serializers;

(c) said switching control further comprises issuance by said controller of a global change-connection-memory-pages signal to every one of said data switches and to every one of said data serializers;

(d) upon receipt of said global change-connection-memory-pages signal, each one of said data switches delays said

switching control between said respective active and standby switch control memories for a time interval corresponding to said start-of-frame time signal contained in said start-of-frame delay register of said one of said data switches; and,

(e) upon receipt of said global change-connection-memory-pages signal, each one of said data serializers delays said switching control between said respective active and standby switch control memories for a time interval corresponding to said start-of-frame time signal contained in said start-of-frame delay register of said one of said data serializers.

4. A switch fabric as defined in claim 1, wherein said egress ports of each one of said data switches in one of said stages are connected to said ingress ports of each one of said data switches in a sequentially next one of said stages in accordance with the equation:

$$\langle p, i, d, pn, egress\rangle :: \langle p, i+1, pn \ div \ portsPerTSE, (d*portsPerTSE) + (pn \bmod portsPerTSE), ingress\rangle$$

where $portsPerTSE = TSEports \ div \ depth$
for p in $0 \ldots planes - 1$
for i in $0 \ldots stages - 1$
for d in $0 \ldots depth - 1$
for pn in $0 \ldots TSEports - 1$

wherein *egress* is an output port address of said data switch, *ingress* is an input port address of said data switch, i is an integer representative of one of said stages, pn is an integer representative

of a pair comprising one of said data switch ingress ports and one of said data switch egress ports, *TSEports* is an integer representative of the total number of said data switch ingress and egress ports for each one of said data switches, *planes* is an integer representative of the total number of said planes, *stages* is an integer representative of the total number of said stages, *depth* is an integer representative of the total number of said depths, and *portsPerTSE* is an integer representative of the total number of egress ports in every one of said data switches within a first one of said stages connected to each one of said data switches within a sequentially next one of said stages.

5. A switch fabric as defined in claim 4, wherein said data serializer egress ports are connected to said data switch ingress ports in accordance with the equation:

$$\langle a, e, egress\rangle :: \langle e \text{ div } portsPerPlane,\ 0,\ (a * portsPerPlane) \text{ div } TSEports,\ ((a * portsPerPlane) \bmod TSEports) + e \bmod portsPerPlane,\ ingress\rangle$$

where $portsPerPlane = TBSports \text{ div } planes$
for a in $0 \ldots TBScount - 1$
for e in $0 \ldots TBSports - 1$

wherein *portsPerPlane* is the number of said data serializer ports in each one of said planes, *TBSports* is the number of said data serializer ingress and egress port pairs for each one of said data serializers, *TBScount* is the total number of said data serializers, *a* is an integer representative of an address of one of said data serializers, and *e* is an integer representative of a port on said one of said data serializers.

6. A switch fabric as defined in claim 5, wherein said data switch egress ports are connected to said data serializer ingress ports in accordance with the equation:

$$\begin{aligned}&\langle p, stages-1, d, pn, egress\rangle ::\\ &\quad \langle((d*TSEports)+pn)\ \text{div}\ portsPerPlane, pn\ \text{mod}\ portsPerPlane\\ &\quad +p*portsPerPlane, ingress\rangle\\ &\\ &\quad \text{where}\ portsPerPlane = TBSports\ \text{div}\ planes\\ &\quad \text{for}\ p\ \text{in}\ planes-1\\ &\quad \text{for}\ d\ \text{in}\ 0\ldots depth-1\\ &\quad \text{for}\ pn\ \text{in}\ 0\ldots TSEports-1\end{aligned}$$

7. A switch fabric as defined in claim 1, wherein:
   (a) each one of said data switches further comprises two input-output connection pins for each pair of said ingress and egress ports of said one of said data switches; and,
   (b) each one of said data serializers further comprises two input-output connection pins for each pair of said ingress and egress ports of said one of said data serializers.

8. A switch fabric as defined in claim 1, wherein said high speed serial links are low voltage differential signalling links.

9. A method of synchronizing operation of a switch fabric as defined in claim 1 for transport of plesiochronous signals by said switch fabric, said method comprising:
   (a) allocating a pre-selected code character to mark a frame boundary beginning;

(b) initializing a frame counter upon detection of a first occurrence of said pre-selected code character in a data stream input to said switch fabric;

(c) monitoring said data stream to detect subsequent occurrences of said pre-selected code character in said data stream;

(d) re-initializing said frame counter after processing each fame of data contained in data stream; and,

(e) signaling misalignment of a frame within said data stream if said pre-selected code character is not detected in said data stream upon said re-initializing of said frame counter.

10. A method as defined in claim 9, further comprising, after signaling said misalignment, re-aligning frames within said data stream by re-initializing said frame counter upon detection of a next occurrence of said pre-selected code character in said data stream.

11. A method as defined in claim 9, wherein said pre-selected code character is an 8b/10b comma character.

12. A method as defined in claim 10, further comprising:

(a) distributing to each one of said data switches and each one of said data serializers a common clock signal generated by a single clock signal generator;

(b) at each one of said data switches and each one of said data serializers:

(i) processing said clock signal to recover a phase component of said clock signal; and,

(ii) monitoring said data stream to recover alignment of said code characters in said data stream.

13. A method of wiring a switch fabric as defined in claim 6, said method comprising providing a lookup table for translating a physical address of each one of said ports into a logical address.

14. A switch fabric as defined in claim 1, further comprising a first memory switch coupled between said data serializer input bus and said data serializer egress ports and a second memory switch coupled between said data serializer ingress ports and said data serializer output bus, said first memory switch for transmitting data received at any time on said input bus at any one of said data serializer egress ports, said second memory switch for outputting data received at any one of said data serializer egress ports at any time on said output bus.

15. A switch fabric as defined in claim 14, wherein:

(a) each one of said data switches further comprises:

(i) a first active memory for each one of said ingress time switches for containing indicia defining currently active connections between selected ingress time slots;

(ii) a first standby memory for each one of said ingress time switches for containing indicia defining next active connections between said selected ones of said ingress time slots;

(iii) a second active memory for containing indicia defining currently active connections between selected ones of said data switch ingress ports, said data switch egress ports and said space switch;

(iv) a second standby memory for containing indicia defining next active connections between said

selected ones of said data switch ingress ports, said data switch egress ports and said space switch;

(v) a third active memory for each of said egress time switches containing indicia defining currently active connections between selected ones of said egress time slots;

(vi) a third standby memory for containing indicia defining next active connections between said selected ones of said egress time slots;

(b) each one of said data serializers further comprises:

(i) an active memory for containing indicia defining currently active connections between selected ones of said ingress and egress ports for said one of said data serializers; and,

(ii) a standby memory for containing indicia defining next active connections between selected ones of said ingress and egress ports for said one of said data serializers;

said switch fabric further comprising a controller for switching control of each one of said data switches and each one of said data serializers between said active and standby memories of said respective data switches and said data serializers.

16. A switch fabric as defined in claim 15, wherein:

(a) each one of said data switches further comprises a start-of-frame delay register for containing a start-of-frame time signal unique to said one of said data switches;

(b) each one of said data serializers further comprises a start-of-frame delay register for containing a start-of-frame time signal unique to said one of said data serializers;

(c) said switching control further comprises issuance by said controller of a global change-connection-memory-pages signal to every one of said data switches and to every one of said data serializers;

(d) upon receipt of said global change-connection-memory-pages signal, each one of said data switches delays said switching control between said respective active and standby memories for a time interval corresponding to said start-of-frame time signal contained in said start-of-frame delay register of said one of said data switches; and,

(e) upon receipt of said global change-connection-memory-pages signal, each one of said data serializers delays said switching control between said respective active and standby memories for a time interval corresponding to said start-of-frame time signal contained in said start-of-frame delay register of said one of said data serializers.

17. A switch fabric as defined in claim 14, wherein:

(a) said egress ports of each one of said data switches in one of said stages are connected to said ingress ports of each one of said data switches in a sequentially next of said stages in accordance with the equation:

$$\langle p, i, d, pn, egress\rangle :: \langle p, i+1, pn \ div \ portsPerTSE, (d*portsPerTSE) + (pn \bmod portsPerTSE), ingress\rangle$$

where $portsPerTSE = TSEports \ div \ depth$
for p in $0 \ldots planes - 1$
for i in $0 \ldots stages - 1$
for d in $0 \ldots depth - 1$
for pn in $0 \ldots TSEports - 1$

wherein *egress* is an output port address of said data switch, *ingress* is an input port address of said data switch, *i* is an integer representative of one of said stages, *pn* is an integer representative of a pair comprising one of said data switch ingress ports and one of said data switch egress ports, *TSEports* is an integer representative of the total number of said data switch ingress and egress ports for each one of said data switches, *planes* is an integer representative of the total number of said planes, *stages* is an integer representative of the total number of said stages, *depth* is an integer representative of the total number of said depths, and *portsPerTSE* is an integer representative of the total number of egress ports in every one of said data switches within a first one of said stages connected to each one of said data switches within a sequentially next one of said stages;

(b) said egress ports of each one of said data serializers are connected to said ingress ports of a first stage of said data switches with the equation:

$$\begin{aligned} &\langle a, e, egress\rangle :: \\ &\quad \langle e \ div \ portsPerPlane,\ 0,\ (a * portsPerPlane) \ div \ TSEports,\ ((a * portsPerPlane) \bmod TSEports + e \bmod portsPerPlane,\ ingress\rangle \\ &\quad \text{where } portsPerPlane = TBSports \ \text{div} \ planes \\ &\quad \text{for } a \text{ in } 0 \ldots TBScount - 1 \\ &\quad \text{for } e \text{ in } 0 \ldots TBSports - 1 \end{aligned}$$

wherein *portsPerPlane* is the number of said data serializer ports in each one of said planes, *TBSports* is the number of said data serializer ingress and egress port pairs for each one of said data serializers, *TBScount* is the total number of said data serializers, *a*

is an integer representative of an address of one of said data serializers, and *e* is an integer representative of a port on said one of said data serializers;

(c) said egress ports of a last stage of said data switches are connected to said ingress ports of each one of said data serializers in accordance with the equation:

$$<p, stages - 1, d, pn, egress> :: <((d*TSEports) + pn) \text{ div } portsPerPlane, pn \text{ mod } portsPerPlane + p * portsPerPlane, ingress>$$

where $portsPerPlane = TBSports \text{ div } planes$
for p in $planes - 1$
for d in $0 \ldots depth - 1$
for pn in $0 \ldots TSEports - 1$

18. A switch fabric as defined in claim 14, wherein:

(a) each one of said data switches further comprises two input-output connection pins for each pair of said ingress and egress ports of said one of said data switches; and,

(b) each one of said data serializers further comprises two input-output connection pins for each pair of said ingress and egress ports of said one of said data serializers.

19. A switch fabric as defined in claim 14, wherein said high speed serial links are low voltage differential signalling links.

20. A method of synchronizing operation of a switch fabric as defined in claim 14 for transport of plesiochronous signals by said switch fabric, said method comprising:

(a) allocating a pre-selected code character to mark a frame boundary beginning;

(b) initializing a frame counter upon detection of a first occurrence of said pre-selected code character in a data stream input to said switch fabric;

(c) monitoring said data stream to detect subsequent occurrences of said pre-selected code character in said data stream;

(d) re-initializing said frame counter after processing each fame of data contained in data stream; and,

(e) signaling misalignment of a frame within said data stream if said pre-selected code character is not detected in said data stream upon said re-initializing of said frame counter.

21. A method as defined in claim 20, further comprising, after signaling said misalignment, re-aligning frames within said data stream by re-initializing said frame counter upon detection of a next occurrence of said pre-selected code character in said data stream.

22. A method as defined in claim 20, wherein said pre-selected code character is an 8b/10b comma character.

23. A method as defined in claim 22, further comprising:

(a) distributing to each one of said data switches and each one of said data serializers a common clock signal generated by a single clock signal generator;

(b) at each one of said data switches and each one of said data serializers:

(i) processing said clock signal to recover a phase component of said clock signal; and,

(ii) monitoring said data stream to recover alignment of said code characters in said data stream.

24. A switch fabric as defined in claim 1, wherein:

(i) said ingress ports are grouped to form a plurality of ingress port groups, each one of said ingress port groups comprising two or more of said ingress ports;

(ii) said egress ports are grouped to form a plurality of egress port groups, each one of said egress port groups comprising two or more of said egress ports;

said switch fabric further comprising:

(a) an ingress memory switch coupled to said ingress port groups for selectably temporally rearranging data received at any one of said ingress port groups and for spatially rearranging data received within members of said ingress port groups; and,

(b) an egress memory switch coupled to said egress port groups for selectably temporally rearranging data transmitted via any one of said egress port groups and for spatially rearranging data transmitted within members of said egress port groups.

25. A switch fabric as defined in claim 24, wherein:

(a) each one of said data switches further comprises:

(i) a first active memory for each one of said ingress memory switches for containing indicia defining currently active ingress time slots for each one of said ingress ports within each one of said ingress port groups;

(ii) a first standby memory for each one of said ingress time switches for containing indicia defining next

active ingress time slots for each one of said ingress ports within each one of said ingress port groups;

(iii) a second active memory for containing indicia defining currently active connections between selected ones of said data switch ingress ports, said data switch egress ports and said space switch;

(iv) a second standby memory for containing indicia defining next active connections between said selected ones of said data switch ingress ports, said data switch egress ports and said space switch;

(v) a third active memory for each of said egress memory switches for containing indicia defining currently active egress time slots for each one of said egress ports within each one of said egress port groups;

(vi) a third standby memory for containing indicia defining next active egress time slots for each one of said egress ports within each one of said egress port groups;

(b) each one of said data serializers further comprises:

(i) an active memory for containing indicia defining currently active connections between selected ones of said ingress and egress ports for said one of said data serializers; and,

(ii) a standby memory for containing indicia defining next active connections between selected ones of said ingress and egress ports for said one of said data serializers;

said switch fabric further comprising a controller for switching control of each one of said data switches and each one of said data serializers between said active and standby memories of said respective data switches and said data serializers.

26. A switch fabric as defined in claim 25, wherein:

(a) each one of said data switches further comprises a start-of-frame delay register for containing a start-of-frame time signal unique to said one of said data switches;

(b) each one of said data serializers further comprises a start-of-frame delay register for containing a start-of-frame time signal unique to said one of said data serializers;

(c) said switching control further comprises issuance by said controller of a global change-connection-memory-pages signal to every one of said data switches and to every one of said data serializers;

(d) upon receipt of said global change-connection-memory-pages signal, each one of said data switches delays said switching control between said respective active and standby memories for a time interval corresponding to said start-of-frame time signal contained in said start-of-frame delay register of said one of said data switches; and,

(e) upon receipt of said global change-connection-memory-pages signal, each one of said data serializers delays said switching control between said respective active and standby memories for a time interval corresponding to said start-of-frame time signal contained in said start-of-frame delay register of said one of said data serializers.

27. A switch fabric as defined in claim 24, wherein:

(a) said egress ports of each one of said data switches in one of said stages are connected to said ingress ports of each one of said data switches in a sequentially next of said stages in accordance with the equation:

<p, i, d, pn, egress> ::
<p, i+1, pn div portsPerTSE, (d*portsPerTSE) + (pn mod portsPerTSE), ingress>

where *portsPerTSE* = *TSEports div depth*
for *p* in 0...*planes* – 1
for *i* in 0...*stages* – 1
for *d* in 0...*depth* – 1
for *pn* in 0...*TSEports* – 1

wherein *egress* is an output port address of said data switch, *ingress* is an input port address of said data switch, *i* is an integer representative of one of said stages, *pn* is an integer representative of a pair comprising one of said data switch ingress ports and one of said data switch egress ports, *TSEports* is an integer representative of the total number of said data switch ingress and egress ports for each one of said data switches, *planes* is an integer representative of the total number of said planes, *stages* is an integer representative of the total number of said stages, *depth* is an integer representative of the total number of said depths, and *portsPerTSE* is an integer representative of the total number of egress ports in every one of said data switches within a first one of said stages connected to each one of said data switches within a sequentially next one of said stages;

(b) said egress ports of each one of said data serializers are connected to said ingress ports of a first stage of said data switches with the equation:

$$<a, e, egress> ::$$
$$<e \text{ div } portsPerPlane, 0, (a*portsPerPlane) \text{ div } TSEports, ((a * portsPerPlane) \text{ mod } TSEports + e \text{ mod } portsPerPlane, ingress>$$

where $portsPerPlane = TBSports \text{ div } planes$
for a in $0 \ldots TBScount - 1$
for e in $0 \ldots TBSports - 1$

wherein *portsPerPlane* is the number of said data serializer ports in each one of said planes, *TBSports* is the number of said data serializer ingress and egress port pairs for each one of said data serializers, *TBScount* is the total number of said data serializers, *a* is an integer representative of an address of one of said data serializers, and *e* is an integer representative of a port on said one of said data serializers;

(c) said egress ports of a last stage of said data switches are connected to said ingress ports of each one of said data serializers in accordance with the equation:

$$<p, stages - 1, d, pn, egress> ::$$
$$<((d*TSEports) + pn) \text{ div } portsPerPlane, pn \text{ mod } portsPerPlane + p * portsPerPlane, ingress>$$

where $portsPerPlane = TBSports \text{ div } planes$
for p in $planes - 1$
for d in $0 \ldots depth - 1$
for pn in $0 \ldots TSEports - 1$

28. A switch fabric as defined in claim 24, wherein said high speed serial links are low voltage differential signalling links.

29. A switch fabric as defined in claim 24, wherein:

(a) each one of said data switches further comprises two input-output connection pins for each pair of said ingress and egress ports of said one of said data switches; and,

(b) each one of said data serializers further comprises two input-output connection pins for each pair of said ingress and egress ports of said one of said data serializers.

30. A method of synchronizing operation of a switch fabric as defined in claim 24 for transport of plesiochronous signals by said switch fabric, said method comprising:

(a) allocating a pre-selected code character to mark a frame boundary beginning;

(b) initializing a frame counter upon detection of a first occurrence of said pre-selected code character in a data stream input to said switch fabric;

(c) monitoring said data stream to detect subsequent occurrences of said pre-selected code character in said data stream;

(d) re-initializing said frame counter after processing each fame of data contained in data stream; and,

(e) signaling misalignment of a frame within said data stream if said pre-selected code character is not detected in said data stream upon said re-initializing of said frame counter.

31. A method as defined in claim 30, further comprising, after signaling said misalignment, re-aligning frames within said data stream by re-initializing said frame counter upon detection of a next occurrence of said pre-selected code character in said data stream.

32. A method as defined in claim 30, wherein said pre-selected code character is an 8b/10b comma character.

33. A method as defined in claim 31, further comprising:

(a) distributing to each one of said data switches and each one of said data serializers a common clock signal generated by a single clock signal generator;

(b) at each one of said data switches and each one of said data serializers:

(i) processing said clock signal to recover a phase component of said clock signal; and,

(ii) monitoring said data stream to recover alignment of said code characters in said data stream.